arXiv:1703.00535v2 [stat.ML] 3 Mar 2017

Human Interaction with Recommendation Systems: On Bias and Exploration

Sven Schmit, Carlos Riquelme
Stanford University
{schmit, rikel}@stanford.edu

March 6, 2017

Abstract

Recommendation systems rely on historical user data to provide suggestions. We propose an explicit and simple model for the interaction between users and recommendations provided by a platform, and relate this model to the multi-armed bandit literature. First, we show that this interaction leads to a bias in naive estimators due to selection effects. This bias leads to suboptimal outcomes, which we quantify in terms of linear regret. We end the first part by discussing ways to obtain unbiased estimates. The second part of this work considers exploration of alternatives. We show that although agents are myopic, agents' heterogeneous preferences ensure that recommendation systems 'learn' about all alternatives without explicitly incentivizing this exploration. This work provides new and practical insights relevant to a wide range of systems designed to help users make better decisions.

1 Introduction

We find ourselves surrounded by recommendations that help us make better decisions. However, relatively little work has been devoted to the understanding of the dynamics of such systems caused by the interaction of users. This work aims to understand these dynamics that arise when users combine the recommendations with their own preference when making a decision.

For example, a user of Netflix uses their recommendations to inform what movie to watch. However, this user also has her own beliefs about movies, e.g. based on artwork, synopsis, actors, recommendations by friends, etc. The user thus combines the suggestions from Netflix with her own preferences to decide what movie to watch. Netflix captures data on the outcome, for example by soliciting a rating, to improve its recommendations to the next user. Of course, this pattern is not unique to Netflix, but observed more broadly; across all platforms that use recommendations.

We are interested in gaining a fundamental understanding of the process of combining suggestions from a recommendation system with preferences of individual users. In particular, there are two problems this work addresses. First, we discuss selection bias. Every outcome of a user interacting with an item observed by a recommendation system is, almost tautologically, observed *after* a user has selected the item. However, this data is used to provide recommendations to users *before* they have made a selection. What kind of problems can occur due to this discrepancy? And what can we do about them? The motivation for these questions comes naturally. Biases in the system induce suboptimal outcomes. And, since we only observe outcomes for actions that are chosen, any observable data hides these biases from us. In particular it is impossible to measure their effects without explicit experiments.

The second problem we want to address is of a different nature. Recommendation engines constantly have to deal with a *cold start* problem, where a new user, or new item, has little to no data. This obviously makes recommendation difficult. Thus, beyond recommending the best items to users, it is also vital to learn efficiently about newer items. Sometimes, it is therefore necessary for the recommendation system to

take risks with suggestions in order to get feedback. That is, a recommendation system needs to balance exploitation with exploration. However, exploration is much more complicated than exploitation. Rather than a prediction problem, it is an inference problem that requires us to quantify the uncertainty in our estimates. This problem easily becomes intractable for more sophisticated recommendation algorithms. Besides statistical difficulties, implementing exploration strategies is generally more challenging in practice, because agents have no incentives to aid the system with exploration. However, one may wonder whether diversity of user preferences leads to natural exploration, such that recommendation systems can explicitly focus on exploitation. The second part of this work addresses this problem and shows that there is indeed natural exploration.

All in all, the focus on this work is on models that describe the relation between system and user, rather than the statistical models that learn from data.

1.1 Main results

This paper provides a dynamical model that captures the dynamics of users with heterogeneous preferences, while abstracting away the specifics of recommendation algorithms. In the first part of this work, we show that there is a severe selection bias problem that leads to linear regret. While we argue this selection bias is difficult to address, we also show how one can reframe the question that solicits feedback that circumvents the bias issue. Second, we show that when the algorithm uses unbiased estimates for arms, 'free' exploration occurs and we recover the familiar logarithmic regret bound. This is important because inducing agents to explore is difficult from both a statistical point and strategic point of view.

1.2 Related work

This work roughly intersects with three separate fields of study. *Recommendation systems* Adomavicius and Tuzhilin [2005] in general have attracted much attention. Collaborative filtering, and matrix factorization techniques haven been particularly well studied. There has been less work explicitly on selection bias, though first demonstrated to exist by Marlin [2003], and more evidence provided in Amatriain et al. [2009]. In follow-up work, Marlin et al. [2007] discuss the not-missing-at-random assumption further, and Steck [2010] also deals with this issue. A key assumption in these works is that there is a covariance shift; the distribution of observed ratings is not altered by conditioning on the selection event, but five star ratings are more likely to be observed. More recently, Schnabel et al. [2016] and Joachims et al. [2017] link the selection bias to recent advances in causal inference. On the other hand, the model in this paper does not make the covariance shift assumption. Furthermore, we focus more on modeling user behavior directly, as opposed to posing assumptions on the data generation mechanism.

The different approach of this work is reminiscent of the work on social learning Chamley [2004]. In social learning, agents learn about the state of the world through combining their private information with observations of actions (but not necessarily outcomes) of others. The seminal work of Smith and Sørensen [2000] shows that people do not necessarily converge to the optimum outcome. The closest paper that relates social learning to this work, is Ifrach et al. [2014], and our model of the dynamics is partly inspired by theirs. They discuss how consumer reviews converge on the quality of a product, given diversity of preferences under a reasonable price assumption. We are fundamentally interested in a different question: how people choose among different products and how recommendation systems cope with selection bias.

Finally, we can relate our work on exploration to the multi-armed bandit literature Bubeck and Cesa-Bianchi [2012]. Classically, much attention has been paid to finding algorithms with low regret in a myriad of settings Lai and Robbins [1985], Auer et al. [2002a,b]. However, multi-armed bandits have also been used to understand how human agents interact with systems: for example, how a system can optimally induce myopic agents to explore by using payments Frazier et al. [2014] or by the way the system disseminates information Papanastasiou et al. [2014], Mansour et al. [2015, 2016]. Similar to those works, we use the regret framework to analyse a system with interacting agents. However, the main difference is that in those works, the (myopic) agents need to be incentivized to explore, while in this work we argue that due to the heterogeneous population, alternatives are explored by myopic agents without the need for incentives from

the system. Thus, our focus is not on the incentive issues. There has also been work on 'free exploration' in auction environments, where Hummel and McAfee [2014] show that presence of noise leads to efficient exploration without the need to incorporate exploration into learning systems. In spirit, that work relates closely to ours, though it focuses on a different application and uses a completely different model.

1.3 Organization

In the next section, we introduce a simple model that captures all the ingredients we require to address both the selection bias and exploration questions. In Section 3 we focus on the issue of selection bias. In Section 4, we discuss the results on free exploration. Before concluding, we illustrate our results using simulations in Section 5.

2 Modeling Human-Algorithm interaction

In this section, we propose a model for the interaction between the recommendation system (*server*) and users (*agents*). Each user selects one of the items the server recommends. Crucially, each item (restaurant, hotel, movie etc.) has an intrinsic quality, that is not known a priori to the user nor the server. Furthermore, each user has their own preferences, such that items with similar qualities can be perceived as very different by users. For example, one traveler prefers a hotel on the waterfront, while another prefers a hotel downtown, and yet a third prefers staying close to the convention center. While these hypothetical hotels have the same quality, the value for users differs, and we posit that these users know about their own preferences. However, the general quality of the hotels is not known to the users.

More formally we assume there are K items, labeled $i = 1, \ldots, K$, and each item has a distinct, fixed but unknown, quality $Q_i \in \mathbb{R}$. This aspect models the vertical differentiation between items and, loosely speaking, it will be the task of the server to estimate these qualities. For notational convenience, we will assume $Q_1 > Q_2 > \ldots > Q_K$. At every time step $t = 1, \ldots, T$, a new user arrives. This user selects one of the K items. To do so, the user has a private preference $\theta_{it} \sim F_i$ for each item, drawn from a preference distribution F_i which we make precise later. The value of item i for user t is

$$V_{it} = Q_i + \theta_{it} + \varepsilon_{it} \tag{1}$$

where ε_{it} is additional noise drawn independently from a noise distribution E with mean 0 and finite variance $\sigma_E^2 < \infty$. We see that if the mean of F_i is 0, then the value of users is centered around the quality, but differs among agents. Agents have some idea about what items they prefer, but since the quality is unknown, they cannot optimally select the best item. To aid the agents, the recommendation system provides a recommendation score s_{it}, aggregating the feedback from previous agents. The agent uses her own preferences, along with the score, to select item a_t according to

$$a_t = \arg\max_i s_{it} + \theta_{it}. \tag{2}$$

Hence, we make the assumption that the agent is boundedly rational and uses s_{it} as a surrogate for the quality. We discuss this assumption further below, but first we finish our exposition of the model. Abusing notation, we write

$$V_t = V_{a_t t} = Q_{a_t} + \theta_{a_t t} + \varepsilon_{a_t t} \tag{3}$$

for the value of the chosen item for agent t. After the agent selects item a_t and observes the value V_t, the server queries for feedback W_t from the user. Initially, we assume $W_t = V_t$; for example the server can ask the user to rate the selected item. However, the private preferences θ_t of the agent remain hidden. The server uses this feedback to give recommendations to future users. In particular, we require s_{it} to be measurable with respect to the past feedback, that is $\sigma\{a_\tau, W_\tau : \tau < t\}$.

Lastly, we need some performance measure to guide our analysis, and quantify the dynamics of human interaction under different algorithms. We note that at each step, the server outputs K scores. However,

only feedback about the selected item is received. Due to this partial feedback, we measure the performance of a recommendation system in terms of (pseudo-)regret as follows:

$$R_T = \sum_{t=1}^{T} \max_i (Q_i + \theta_{it}) - (Q_{a_t} + \theta_{a_t t}), \tag{4}$$

which sums the difference between the expected value of the best item[1] and the expected value of the selected item. In particular, if scores $s_{it} \equiv Q_i$ for all t, the regret of such server would be 0, and each user picks the optimal action using equation 2. We stress that this notion of regret serves as a tool, rather than a goal; we are not particularly interested in designing an algorithm that minimizes this notion of regret, but rather use regret to analyse simple algorithms, and gain insights into the dynamics of recommendation systems in general.

2.1 A note on incentives

We note that the agents in our model are boundedly rational. Experimentally, there has been abundant evidence of human behavior that is not rational Camerer [1998], Kahneman [2003]. Simple heuristics of user behavior have been used by others in the social learning community. Examples include learning about technologies from word-of-mouth interactions Ellison and Fudenberg [1993, 1995] and modeling persuasion in social networks Demarzo et al. [2003]. The combination of machine learning and mechanism design with boundedly rational agents is explored in Liu et al. [2015].

From the perspective of our model, incentives of server and agents are quite aligned, except that agents are myopic and the server also cares about future agents. This implies that agents just want to select the item that is best for them, while from the server's standpoint, it might be better to nudge the agent to select a different item in order for the server to provide better recommendations to future agents. For example, the server could boost the score of a particular item to increase the probability it is selected by the agent. However, we are not interested in such algorithms. In fact, we argue in Section (4) that there is no need to do this to obtain order-optimal performance. And if the server supplies point estimates that help agents in their myopic choice, there is no tension between server and agent. In particular, we do note that if the server outputs the true qualities Q_i, then the selection rule (2) is optimal.

2.2 Bernoulli preferences

Because analysis for general distributions F_i over preferences is intractable, we focus on the simpler case where $Q_i \in [0, 1]$ and F_i is Bernoulli distributed with success probability p_i. Furthermore, to avoid complicating notation, we assume $p_i = p$ is the same across all items, though this is not strictly necessary and results can be extended to the generalized case.

When there are many items, it does not make sense to have a large fraction of items, all with $\theta_{it} = 1$, because it leads to vacuous results. If there is always a large fraction of items with a positive signal, the model can be seen as a sleeping bandits problem Kleinberg et al. [2008], where items without such signal are considered sleeping, and the choice is based on the remaining options. Therefore, one should think of $p \lesssim \frac{\log(K)}{2K}$, which ensures that the number of items with a positive signal for a user is approximately constant as we vary K.

For convenience, we sometimes write that user t has observed a *positive signal* for item i if $\theta_{it} = 1$ and, likewise a *negative signal* if $\theta_{it} = 0$. Based on this model, we investigate its properties in the next sections, but first we provide some more remarks regarding the proposed model.

2.3 A note on personalization

It is clear that this model lacks personalization; the quality that the server tries to estimate is the same across users. To add personalization, we can replace the scalar quality Q_i by some function based on covariates.

[1] Unlike the traditional bandit setting, the best item is not necessarily item 1. Rather, it depends on the user preferences.

Then we can use more sophisticated machinery to fit such model. However, the focus of this work is not on statistical learning techniques, but rather on understanding the dynamics of the feedback loop between recommendation system and humans and preferences.

One could argue that better personalization, that is, more complex statistical models which capture additional information such as covariates and the history of the user, is able to model the heterogeneous user preferences. Put in terms of our model, this argument suggests that θ_{it} can be absorbed into s_{it}. We believe there is some truth to that; it can model some of the heterogeneous preferences.

However, we argue that in most, if not all, cases this factor cannot be completely eliminated. There is some part of the user's preferences that cannot be captured by covariates and historical data. This implies that there is an 'unobservable' component to the user's choice, denoted by θ_{it} in our case. There are two strong arguments in favor of this. Both are based on the fact that every recommendation system is constrained in terms of the quantity *and* quality of the data it is based on. From a quantity standpoint, a user only interacts with a system so often, and that limits the amount of personalization that models can achieve. Even if there are models that could capture all of the preferences, in practice systems lack the data to support such models. This implies that some part of the preferences remain beyond the scope of models. Second, not only quantity, but also quality of data is important. Often, recommendation systems have access to only few features, and some aspects of user preferences of user preferences, such as taste or style, can be difficult to capture. These constraints make it difficult to fully model users preferences, validating the need to explicitly model the unobserved preferences to get a deeper understanding of the dynamics of recommendations systems.

3 Biased estimates

In this section, we analyse the performance of 'naive' algorithms, that is, scoring processes that do not take into account that agents have private preferences, and base the scores on empirical averages. This is equivalent to the server assuming $F_i \equiv 0$. We focus on the Bernoulli preferences model, though in Section 5 we empirically demonstrate similar outcomes for different preference distributions.

First we define the set of agents before time t that have selected item i by

$$S_{it} = \{\tau < t : a_\tau = i\}. \tag{5}$$

We also define $\bar{V}_{it}$ to denote the empirical average of item i up to time t:

$$\bar{V}_{it} = \frac{1}{|S_{it}|} \sum_{\tau \in S_{it}} V_\tau \tag{6}$$

where we assume that for $S_{it} = \emptyset$, $\bar{V}_{it} = 0$, for the average observed outcome for the agents that selected item i before time t. We want to show that the system suffers linear regret when the server uses any scoring mechanism for which scores converge to the empirical average of the observed values. To make this rigorous, we define the notion of *mean-converging* scoring process.

Definition 1. *A scoring process that outputs scores s_{it} for item i at time t is* mean-converging *if*

1. *s_{it} is a function of $\{V_\tau : \tau \in S_{it}\}$ and t.*
2. *$s_{it} \to \bar{V}_{it}$ if $\liminf_t \frac{|S_{it}|}{t} > 0$ almost surely.*

That is, the score only depends on the observed outcomes for this particular item, and if we observe a linear number of selections of arm i, then the score converges to the mean outcome. Trivially, this includes using the average itself as score, $s_{it} = \bar{V}_{it}$, but also includes well known methods that carefully balance exploitation with exploration, such as versions of UCB and Thompson Sampling.

Lemma 1. *An upper confidence bound strategy with an upper bound of the form $\alpha\sqrt{\frac{\log(T)}{|S_{it}|}}$ is mean-converging.*

This is immediate from the definition of mean-converging. The same is true for Thompson sampling as long as the prior is independent and well specified.

Lemma 2. *If the noise distribution E_i has a normal distribution, then Thompson sampling with an independent normal prior for each Q_i is mean-converging.*

This follows because the prior washes out and therefore the posterior will converge to the empirical point estimate.

From the previous section, we know that ideally the scores supplied to the user converge to the quality of the item, $s_{it} \to Q_i$, as more users select item i. We say that the scores are *biased* if this is not the case:

$$s_{it} \not\to Q_i \quad \text{as} \quad |S_{it}| \to \infty. \tag{7}$$

The next proposition shows that mean-converging scoring processes lead to linear regret, because these scores are generally biased. However, it does require a 'gap condition' on the differences in quality that is sufficiently small. We illustrate this gap effect with simulations in Section 5.

Proposition 3. *Under Bernoulli preference model with $p_i = p \in (0, 1)$, if*

$$\Delta = Q_1 - Q_2 < \frac{(1-p)^K}{(1-p)^K + p} \tag{8}$$

and s_{it} is mean-converging, then

$$\limsup_{t\to\infty} \frac{R_t}{t} \geq c \tag{9}$$

for some $c > 0$.

The proof of this proposition can be found in the appendix. For $K = 100$ and $p = 1/50$, the condition requires $\Delta < 0.86$. More generally in the relevant regime where $p < \frac{\log(K)}{2K}$, the condition on is satisfied if $\Delta < 0.7$ for all K. We also note that the linear regret we obtain does not have to do with the usual exploration/exploitation trade-off, but rather our estimators being biased. One wonders if the bias is equally problematic under different distributional assumptions on the private preferences. Based on simulations[2] we have found that issues arise generally, though they are more difficult to characterize as they depend on the specific distributions (and parameters).

3.1 Unbiased estimates

Naturally, a first attempt to improve the linear regret is aimed at obtaining unbiased versions of the naive averaging. We sketch two approaches; one based on algorithmically adjusting estimates, and the other attempting to avoid the bias to begin with. The latter could be achieved by changing the type of feedback we request from users.

3.1.1 Algorithmic approach

We first discuss algorithmic solutions. In the case of Bernoulli preferences with a common parameter p, we can estimate p using data on the ranks of chosen items, which follows a truncated negative binomial distribution. Furthermore, we note that the feedback is biased upwards by 1 for any item but the first ranked one. For the item ranked first the feedback is upward biased by $\frac{p}{p+(1-p)^K}$; the probability that it is chosen with a positive signal. Therefore, to unbias, we estimate p from the data we have, and depending on whether the selected item was ranked first or not, we subtract either $\frac{\hat{p}}{\hat{p}+(1-\hat{p})^K}$ or 1, respectively, from the reported value. If we relax the assumption of a common p, but instead consider a Bernoulli model where each item i has a corresponding p_i, these can be estimated from data, and we can basically proceed as before.

[2]For example, by using normal distributions for the preferences.

However, if presented with real data, we would be hesitant to impose such rigid structure on the preferences of users. While the Bernoulli model is helpful in yielding a tractable model that gives insight into the main dynamics, it is not realistic to think it accurately reflects the preferences in practice. At a minimum, it makes sense to use a continuous preference distribution. However, algorithmically unbiasing the data appears impracticable and intractable in this setting; impracticable because it still requires strong assumptions on the preference distributions, and intractable because it requires integrating over the preference distributions for each item.

3.1.2 Changing the feedback model

Instead, we prefer to look for an alternative. And as hinted at before, one way is to request different feedback from the user. The traditional type of question 'How would you rate this item?' asks for an absolute measure of satisfaction, which corresponds to directly probing for V_t in our model. But, this measurement is flawed because both the choice and the value are driven by the preference vector θ_t. To avoid this problem, one can imagine that initially the user expects a value of $s_{it} + \theta_{it}$ for item i. If we ask how the chosen item thus compared to the expectation, we ask for a relative measure of feedback, approximating $V_t - (s_{a_t t} + \theta_{a_t t})$. An example of such prompt could be 'How does this item compare to your expectation?'. Given that the server stores the supplied score $s_{a_t t}$, we note that we can uncover an unbiased estimate of Q_{a_t}. This avoids having to deal with the bias at all, and, importantly, it does not require any distributional assumptions on the form of the preferences.

Asking for relative feedback thus seems appealing, but also comes with two caveats. The more obvious one is whether users are actually able to give such relative feedback in a reliable way. One can argue that users might already give such feedback implicitly, or that 'a priori expectation' is something users are not capable of accurately reporting. The second caveat comes from the fact that many recommendation systems now rely more on implicit feedback, e.g., did someone finish watching the movie. Such measures are necessarily absolute and our suggestion is vacuous in such scenarios.

Concluding this section and the first part of our results, we have discussed two alternatives that try to obtain unbiased estimates. In particular, changing the way recommendation systems ask for feedback could be a viable option to reduce selection bias.

4 Exploration

From the previous section we know that biased scoring mechanisms lead to linear regret. In this section, we investigate what happens with simple but unbiased scoring rules. Thus, we assume we have access to unbiased feedback from now on, obtained either from users or algorithmically. But, this is not necessarily sufficient to guarantee good performance. From the multi-armed bandit literature, we know that for a system to have low regret, it needs to carefully balance exploration with exploitation.

In our case, that means that the system needs to obtain data on every item in order to provide useful scores to the users. However, we cannot expect users to be so kind to do 'dirty work' for the server. Therefore, there has recently been an increased interest in understanding this trade-off in the presence of myopic agents, who are naturally interested in doing well for themselves, rather than helping the server learn Frazier et al. [2014], Mansour et al. [2015, 2016], Papanastasiou et al. [2014]. The previously mentioned works all address the question of how to incentivize users to not act myopicly in different ways.

In this section we address the problem of exploration in the proposed model. As opposed to the research just mentioned, we deal with agents with heterogeneous preferences. It seems natural that these heterogeneous agents help the system explore, but it is not obvious to what extent this helps. We show that because of this diversity in preferences, the free exploration leads to optimal performance of the system up to constants; we recover the standard logarithmic regret bound from the bandit literature. This means that there is little need for a server to implement a complicated exploration strategy, and incentives naturally align much better than in the settings of previous work.

4.1 Formal result

The formal result presented below again assumes the Bernoulli preferences model. But in addition, we assume access to unbiased feedback from the user. That is, the feedback at time t for chosen item i_t is $\tilde{V}_t = Q_{i_t} + \varepsilon_t$. We consider a server that uses empirical averages of these unbiased scores,

$$\tilde{s}_{it} = \frac{1}{|S_{it}|} \sum_{\tau} \tilde{V}_\tau \tag{10}$$

However, this is not enough to ensure a tight regret bound, as a single dramatically low rating for some item can cause all future agents to ignore that item forever. Therefore, we impose the condition that

$$1 + \min_i \tilde{s}_{it} > \max_i \tilde{s}_{it} \tag{11}$$

by increasing the lowest scores to satisfy this criterion if needed. In practice, this ensures that for any item and any set of valid scores there is a realization of preferences for a user such that the item is selected. In general, it is impossible to prove regret bounds for the empirical scores with discrete signals and unbounded support for error terms; with some small but positive probability, an item gets such a terrible rating that the signals cannot make a difference. From a practical standpoint, it would also not make sense to offer items for which the server knows no one is interested. Note that this does not lead to incentive issues because we assume all qualities lie in the unit interval.

Under this extra condition, we can prove that empirical averaging is enough to get an order optimal (pseudo)-regret bound with respect to the total number of agents T.

Proposition 4. *Assume ε_t is σ-sub-Gaussian, and the server uses s_{it} as scores. Then*

$$\mathbb{E}[\text{regret}(T)] \leq \frac{32\alpha\sigma^2 K(\log(T) - \log(\Delta_{\min}) + \log(2))}{\Delta_{\min}^2 p(1-p)^K} + \left(\frac{16\sigma^2}{\Delta_{\min}} + \Delta_{\min}\right) K \tag{12}$$

where α is a constant depending on σ^2 and $\Delta_{\min}$.

Some remarks are in order. First, we note that this is a problem dependent bound, in that it depends on the gaps of quality between items. It is more difficult to get around this issue than in the standard setting, because each item is optimal with some probability. Furthermore, bound is not optimized in terms of K, and the constants are large.

The proof can be found in the Appendix A.2. The gist of the proof is first showing that after a certain amount of steps, each arm has been tried a sufficient amount by random exploration of agents. This then implies that the empirical estimates are close to the true qualities with high probability. From this we can then deduce that no further agents choose suboptimal actions. The main take-away from this result is not so much the specific bound, but rather the practical insight that it, and its proof, yield. The intuition is that initially estimates of quality are poor. Therefore, it takes some time and luck for users with idiosyncratic preferences to try these items. As estimates improve, however, most agents are drawn to their optimal choice. Since these choices differ across agents, we now get to learn efficiently without incurring a regret penalty.

In standard bandit problems, it is important to ensure each action is tried a logarithmic number of times to avoid getting stuck with a suboptimal action. Our analysis shows that this is not the case here. The bottleneck really comes from items with too few observations. After a sufficient number of data points, the estimated quality will be good enough in the sense that user diversity kicks in to provide convergence for free. The practical consequence of this observation is that to improve the performance, the designer of a recommendation system should focus on simple ways to make new items, or more generally items with few observations, more likely to be chosen. This can be achieved by highlighting new arrivals, something that is commonly observed in practice. For example, Netflix has a 'Recently Added' selection that is often clearly displayed to the user.

5 Simulations

In this section, we look at simulations of the model that give more insight into the dynamics of the system. First, we visualize the effect of the gap condition in a simple setting with two arms. Second, we demonstrate the effect of free exploration on cumulative regret by running simulations with biased and unbiased scores.[3]

5.1 Demonstration of the gap effect on bias

Proposition 3 requires a condition on the gap between the best and second best arm. This is caused by the discrete and bounded nature of the preferences of users in the Bernoulli preferences model. To improve our understanding of this phenomenon, we simulate two sample paths from a simple model with only 2 arms that makes it easy to visualize the effect.

The quality of the bad arm is set to zero, $Q_2 = 0$, and we vary the quality of the good arm $Q_1 = \Delta$, where we consider $\Delta = 0.2$ and $\Delta = 0.5$. The private preferences are drawn from Bernoulli distributions with probability 0.5 of observing a success. We then run the system as 2000 agents sequentially arrive, get a recommendation score, select an item and leave feedback according to the model we described in Section 2.

In Figure 1 we show the dynamics of the scores s_{it} on the left and the fraction of times each arm is selected, averaged over a local window using exponentially weighted averaging with $\alpha = 0.98$, on the right. The plot in the top left corner shows that when the gap in quality is large, there is also a gap in the two scores s_{1t} and s_{2t}. Therefore agents to always select their optimal action. Note that there is bias in the scores, and in particular the difference between the scores $(s_{1t} - s_{2t})$ is much smaller than the gap in qualities $(Q_1 - Q_2)$. On the top-right we indeed see that the fraction of time each item is selected tracks the optimal rate for both items; only agents that have a negative preference for the first item and a positive preference for the second item choose it, which happens with probability 25%.

On the other hand, when the gap in qualities is small, as shown in the two plots at the bottom, we actually notice that the scores closely track each other. This causes a much larger fraction of users to select the second item, as shown in the bottom-right plot. In particular, the recommendation system is unable to help agents with $\theta_{1t} = \theta_{2t}$ select the first item.

5.2 Demonstration of regret

The goal of the next simulation is to demonstrate the effect of free exploration when estimates are unbiased, and leads to logarithmic as opposed to linear regret. In this case, we run a simulation with 50 items, where each quality is drawn uniformly from the unit interval, and with 5000 agents. This allows us to look at a reasonably sized problem, while also being able to visualize the dynamics of regret over a reasonable timescale. Furthermore, the preferences for items are drawn from a Bernoulli distribution with $p = \frac{\log(k)}{2k} \approx 0.04$, in line with the relevant regime mentioned in Section 3.

Figure 1 shows the cumulative regret, and we note that initially, both biased and debiased averages incur high regret. However, as both algorithms get more information, we note that the regret for debiased averages flattens dramatically while the naive averages show a linearly increasing pattern all the way to the 5000th agent.

6 Discussion

In this work, we introduce a model for analyzing feedback in recommendation systems. We propose a simple model that explicitly looks at heterogeneous preferences among users of a recommendation system, and takes the dynamics of learning into account. Two phenomena occur in this model: selection bias and free exploration. Selection bias is caused by the diversity of preferences among users, and it leads to linear regret. We can either algorithmically try to unbias the feedback, or reframe the feedback by asking for relative outcomes rather than absolute outcomes. That is, ask 'How does this compare to your expectation?', rather

[3] The code to replicate the simulations is publicly available at `https://github.com/schmit/human_interaction`

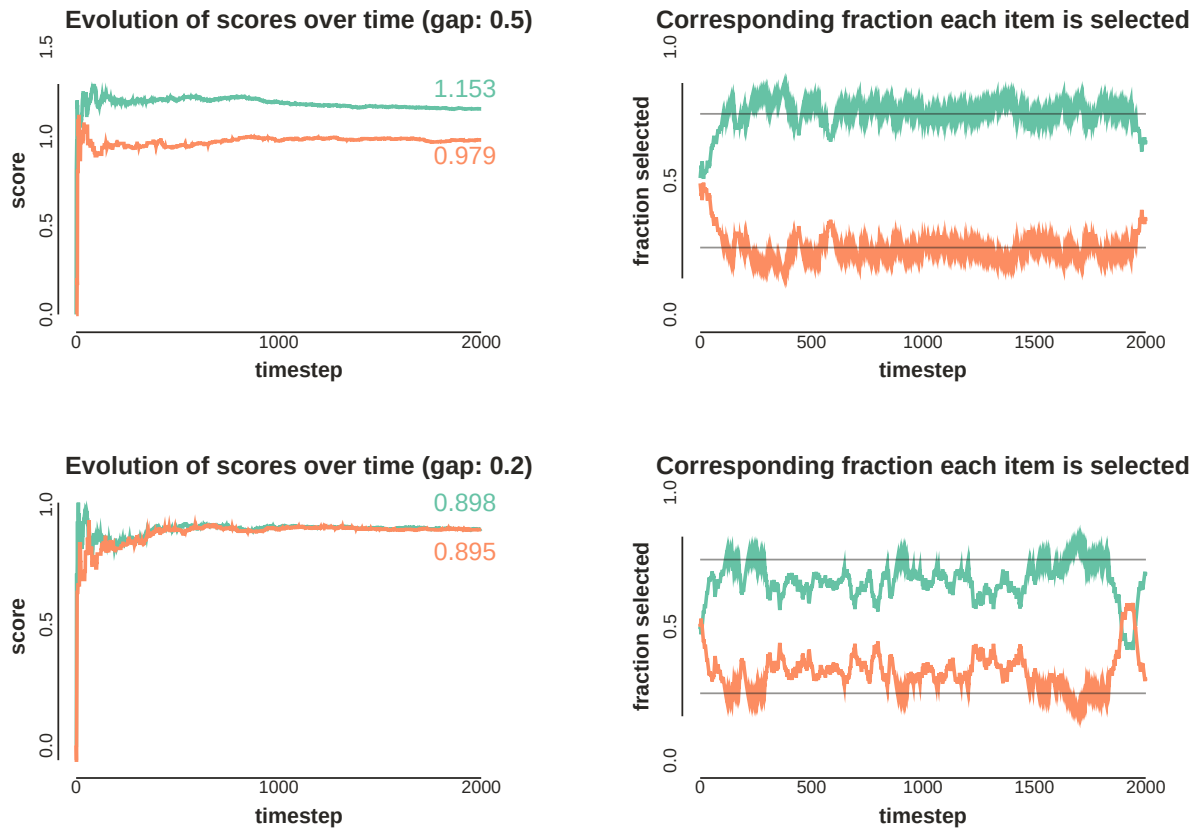

Figure 1: These plots show two simulations. At the top, we simulate a model where the gap between qualities is large, while at the bottom, the gap is small. On the left, the evolution of recommendation scores (s_{it}) is shown, while on the right we show the fraction of times each item is selected, smoothed using exponentially weighted averaging. The green lines correspond to the good arm, while the orange lines correspond to the bad arm. The black lines on the plots to the right indicate the optimal selection rate.

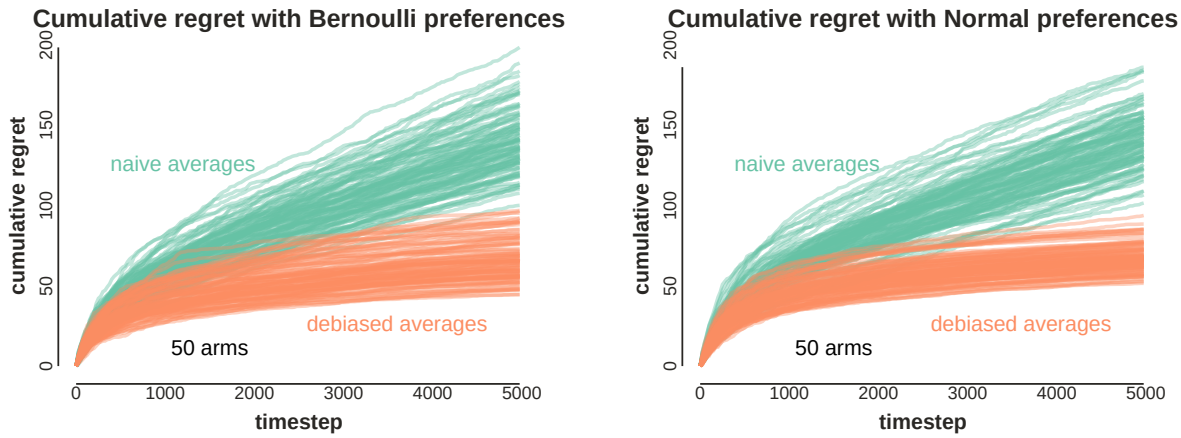

Figure 2: These plots shows the cumulative regret plotted against time for both the naive (biased) empirical averages in green, and the debiased averages in orange. We plot 100 paths for each, where the qualities of 50 arms are drawn uniformly at random from the unit interval. On the left, preferences follow the Bernoulli model, while on the right the preferences are drawn from the standard Normal distribution.

than 'How do you like this?'. Algorithmically adjusting for bias seems unreasonable in practice, as it relies on stringent assumptions. Furthermore, we recognize that reframing feedback is not always straightforward to implement in practice. However, it sheds new light on how one might solicit feedback in recommendation systems.

The second part of this work discusses the phenomenon of free exploration. That is, the diversity of preferences ensures that the system learns about each item despite not explicitly forcing exploration. Because of this, the server continues to learn about every item, and we recover the standard logarithmic regret rate. This stands in sharp contrast to other work on exploration with myopic agents, where explicit and complicated incentives are required. From a practical point of view, our analysis shows that the common sense approach of highlighting new items aligns with improving exploration, as only items with limited feedback affect regret adversely.

6.1 Future work

There are several directions of further research. In our simple model, there are many aspects of recommendation systems that are not captured. The most interesting aspect is that in practice users only observe a limited set of recommendations, rather than the entire inventory. This is something we have not modeled and has implications on the rate of exploration.

Beyond exploration, things get more interesting when there are features that are correlated with both the feedback (V_t) and the user selection (a_t). To do well while only supplying a limited set of recommendations, the server has to combine an outcome model (based on the rating) with a selection model. It is unclear how to do this optimally, as it requires balancing showing items that users are likely to select with items that they are likely to rate highly. For example, if a certain feature is a strong predictor for selection, it is likely a weak predictor for outcome conditioned on selection; if users base their selection heavily on a particular feature, then there is little variance left to exploit for the feedback part of the model.[4]

Binary feedback is another avenue for further work. On one hand this coarser feedback reduces the

[4] This is best illustrated with an example. Suppose users base their movie selection on genre, such that a user that loves comedies only selects comedies, then there is no feedback for thrillers. Therefore, the server has trouble picking up on the correlation between genre and feedback, even though this effect could be strong.

efficiency in learning, but on the other hand we believe that this feedback mechanism is more robust to changes in modeling assumptions and more common in practice. Moreover, it would be interesting to develop additional theory and run simulations with latent variable models and collaborative filtering techniques such as matrix factorization.

6.2 The bigger picture

Beyond the model we propose and its analysis, which leaves plenty of questions unanswered, we believe that this work has raised fundamental and important issues relating the interaction between machine learning systems and the users interacting with them. Algorithms not only consume data, but in their interaction also create data, a much more opaque process but equally vital in designing systems that achieve the goals we set out to achieve. It is therefore important to not only improve state-of-the-art algorithms, but also improve our understanding of the input to those algorithms.

7 Acknowledgements

The authors would like to thank Ramesh Johari, Vijay Kamble, Brad Klingenberg, Yonatan Gur and Peter Lofgren for their suggestions and feedback. This work is supported by the National Science Foundation.

References

G. Adomavicius and A. Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Trans. Knowl. Data Eng.*, 17:734–749, 2005.

X. Amatriain, J. M. Pujol, and N. Oliver. I like it... i like it not: Evaluating user ratings noise in recommender systems. In *International Conference on User Modeling, Adaptation, and Personalization*, pages 247–258. Springer, 2009.

P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, 2002a.

P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM J. Comput.*, 32:48–77, 2002b.

S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *CoRR*, abs/1204.5721, 2012.

C. Camerer. Bounded rationality in individual decision making. *Experimental economics*, 1(2):163–183, 1998.

C. Chamley. *Rational Herds: Economic Models of Social Learning*. Rational Herds: Economic Models of Social Learning. Cambridge University Press, 2004. ISBN 9780521530927. URL `https://books.google.com/books?id=2dgbOh6VE9YC`.

P. M. Demarzo, D. Vayanos, J. Zwiebel, N. Barberis, G. Becker, J. Bendor, L. Blume, S. Board, E. Dekel, S. Dellavigna, D. Duffie, D. Easley, G. Ellison, S. Gervais, E. Glaeser, K. Judd, D. Kreps, E. Lazear, G. Loewenstein, L. Nelson, A. Neuberger, M. Rabin, J. Scheinkman, A. Schoar, P. Sorenson, P. Veronesi, and R. Zeckhauser. Persuasion bias, social influence, and unidimensional opinions. 2003.

G. Ellison and D. Fudenberg. Rules of thumb for social learning. *Journal of Political Economy*, 101(4):612–643, 1993.

G. Ellison and D. Fudenberg. Word-of-mouth communication and social learning. *The Quarterly Journal of Economics*, 110(1):93–125, 1995.

P. I. Frazier, D. Kempe, J. M. Kleinberg, and R. Kleinberg. Incentivizing exploration. In *SIGECOM*, 2014.

P. Hummel and R. P. McAfee. Machine learning in an auction environment. In *WWW*, 2014.

B. Ifrach, C. Maglaras, and M. Scarsini. Bayesian social learning with consumer reviews. *SIGMETRICS Performance Evaluation Review*, 41:28, 2014.

T. Joachims, A. Swaminathan, and T. Schnabel. Unbiased learning-to-rank with biased feedback. *CoRR*, abs/1608.04468, 2017.

D. Kahneman. Maps of bounded rationality: Psychology for behavioral economics. *The American economic review*, 93(5):1449–1475, 2003.

R. D. Kleinberg, A. Niculescu-Mizil, and Y. Sharma. Regret bounds for sleeping experts and bandits. In *COLT*, 2008.

T. L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1): 4–22, 1985.

T.-Y. Liu, W. Chen, and T. Qin. Mechanism learning with mechanism induced data. In *AAAI*, pages 4037–4041, 2015.

Y. Mansour, A. Slivkins, and V. Syrgkanis. Bayesian incentive-compatible bandit exploration. *CoRR*, abs/1502.04147, 2015.

Y. Mansour, A. Slivkins, V. Syrgkanis, and Z. S. Wu. Bayesian exploration: Incentivizing exploration in bayesian games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 661–661, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-3936-0. doi: 10.1145/2940716.2940755. URL `http://doi.acm.org/10.1145/2940716.2940755`.

B. M. Marlin. Modeling user rating profiles for collaborative filtering. In *NIPS*, 2003.

B. M. Marlin, R. S. Zemel, S. T. Roweis, and M. Slaney. Collaborative filtering and the missing at random assumption. In *UAI*, 2007.

Y. Papanastasiou, K. Bimpikis, and N. Savva. Crowdsourcing exploration. *History*, 2014.

T. Schnabel, A. Swaminathan, A. Singh, N. Chandak, and T. Joachims. Recommendations as Treatments: Debiasing Learning and Evaluation. *ArXiv e-prints*, Feb. 2016.

L. Smith and P. Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.

H. Steck. Training and testing of recommender systems on data missing not at random. In *KDD*, 2010.

M. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*. Forthcoming, 2015.

A Appendix

A.1 Proof bias

Proof of Proposition 3. Fix a sample path $\omega \in \Omega$. Note that by assumption, each arm is optimal for a constant fraction of agents. Let x_{it} denote the empirical fraction of time arm i has been pulled up to time t. That is,

$$x_{it} = \frac{|\{\tau : a_\tau = i\}|}{t}. \tag{13}$$

Then, if $\liminf_t x_{it} < x_i^*$ for some sufficiently small $x_i^* > 0$, we incur linear regret almost surely. Instead, assume that each arm is sampled a constant fraction, $\liminf_t x_{it} > c_{xi}$ for some c_{xi} for each arm i, which implies that we continue learning about each arm. We note that the expected reward for the item ranked highest is

$$Q + \frac{p}{p + (1-p)^K} = Q + \rho, \tag{14}$$

where Q is the quality of this item and we define $\rho = \frac{p}{p+(1-p)^K}$. With probability p this item is chosen because of a positive signal, and with probability $(1-p)^K$ it is chosen because none of the items have a positive signal.

For the other items, the expected reward is $Q_i + 1$, as they are only chosen when the agent receives a positive signal. Recall that without loss of generality, we assume $Q_1 > Q_2 > \ldots > Q_K$. Furthermore, it is convenient to define c_t as the fraction (up to time t) that the item i is not ranked at the top:

$$c_t = \frac{|\{\tau < t : \exists j > 1 : s_{1\tau} < s_{j\tau}\}|}{t}. \tag{15}$$

We note that if $\limsup_t c_t > c^*$ for some $c^* > 0$, then we have the desired linear regret.

Informally, to show linear regret we want to find

$$\frac{\mathbb{P}(\text{ Select arm } i \text{ while ranked best })}{\mathbb{P}(\text{ Select arm } i\ )}, \tag{16}$$

under any scoring rule such that arm i is ranked first with probability $1-c$. This quantity helps us find the long term average the score for the ith arm will converge to, as it depends on how often this arm is ranked first. However, it is difficult to compute this quantity directly, because it depends on the distribution of ranks for arm 1. We can find simple bounds for this quantity by further conditioning on the rank of arm i when it is selected while it is not ranked highest. The two extreme cases are that it is either ranked second or last in such case. We define λ and λ' as functions of c to denote these conditional probabilities that bound the quantity above.

$$\lambda(c) = \frac{(1-c)(p+(1-p)^K)}{(1-c)(p+(1-p)^K) + cp(1-p)} \tag{17}$$

which denotes the expected fraction an item is selected under the top score given that if it is selected but is not ranked first, then it is ranked second, and

$$\lambda'(c) = \frac{(1-c)(p+(1-p)^K)}{(1-c)(p+(1-p)^K) + cp(1-p)^{K-1}} \tag{18}$$

which denotes the expected fraction selected as the top ranked item given that if the item is selected but does not have highest score, then it was ranked last. Hence, $\lambda(c)$ and $\lambda'(c)$ form two extremes, and it follows that the expected fraction selected under the top score is in between those two values. Note that λ and λ' are both decreasing in c and $\lambda'(c) > \lambda(c)$ for all c.[5]

Now suppose $\limsup c_t = c$. By the stong law of large numbers, the empirical average converges to its mean and thus

$$\limsup_t s_{1t} \leq Q_1 + \lambda(1-c)\rho + (1 - \lambda(1-c)), \tag{19}$$

where the second term corresponds to the expected reward from being ranked first and the last term corresponds to the contribution from when the action is not ranked first. Similarly

$$\liminf_t s_{2t} \geq Q_2 + \lambda'(c)\rho + (1 - \lambda'(1-c)), \tag{20}$$

almost surely by the mean-converging condition.

We note for $c = 0$, this leads to

$$\limsup_t s_{1t} \leq Q_1 + \rho \tag{21}$$

and

$$\liminf_t s_{2t} \geq Q_2 + 1 \tag{22}$$

[5] Both have the form $\frac{(1-x)a}{(1-x)a+xb}$ for $a, b \in (0,1)^2$, which has a negative derivative for $x \in (0,1)$

This is a contradiction if $\Delta < \frac{(1-p)^K}{p+(1-p)^K}$, as this would imply the score of the second arm is higher in the limit than that of the first arm, while the first item is always ranked before the second item ($c = 0$):

$$\limsup_t s_{1t} \leq Q_1 + \rho = Q_2 + \Delta + \rho < Q_2 + \frac{(1-p)^K}{p+(1-p)^K} + \frac{p}{p+(1-p)^K} = Q_2 + 1 \leq \liminf_t s_{2t}. \tag{23}$$

Furthermore, since λ and λ' are continuous and monotone, there exists some $c^* \in (0, 1)$ such that

$$\liminf_t s_{2t} \geq Q_2 + q - \lambda'(c^*)(1-\rho) > Q_1 + 1 - \lambda(1-c)(1-\rho) \geq \limsup_t s_{2t} \tag{24}$$

almost surely. This implies that $\limsup_t c_t > c^*$ almost surely, which proves the linear regret almost surely bound from the proposition. □

A.2 Proof exploration

Proof Proposition 4. We look at individual arms and note that if at time t all estimates are accurate to within λ at such time the regret is at most 2λ. Furthermore, if $\lambda < \frac{\Delta_{\min}}{2}$, the regret is 0, because agents select the optimal arm no matter their preferences. Also note that in general, the regret at any time is at most 2. Hence, the strategy is to find a proper concentration bound on the estimation error after observing at least m selections for that item for some suitable choice of m. We can also find a concentration bound on the time it takes to see m selections of an item, since the probability of observing an item is at least $p(1-p)^K$.

Define events

$$A_m(i, \lambda) = \left\{ \exists s \in \{m, \dots, T\} : \frac{1}{s} \left| \sum_{i=1}^{s} \varepsilon_i \right| > \lambda \right\} \tag{25}$$

and

$$B_m(i, M) = \{|S| < m : \tau \in S \iff a_\tau = i \text{ and } \tau < M\}. \tag{26}$$

That is, $A_m(i, \lambda)$ is the event that after more than m pulls there is a t where the estimate of the quality for arm i is off by more than λ, and $B_m(i, M)$ is the event that we do not observe item i at least m times within M pulls. We set m and M to suitable values later. Based on these two events, we can bound the expected regret by

$$\mathbb{E}[\text{regret}(T)] \leq \sum_{i=1}^{K} 2(\mathbb{P}(B_m(i, M)) + \mathbb{P}(A_m(i, \lambda)))T + 2M + \lambda T \mathbb{I}_{\lambda > \frac{\Delta_{\min}}{2}}. \tag{27}$$

Bounding A_m Using the standard σ-sub-Gaussian concentration bound (see, for example, [Wainwright, 2015, Chapter 2]), we have

$$\begin{aligned}
\mathbb{P}(A_m(i, \lambda)) &\leq \mathbb{P}\left(\exists s \in \{m, \dots, t\} : \frac{1}{s} \left| \sum_{i=1}^{s} \varepsilon_i \right| > \lambda \right) && (28)\\
&\leq \sum_{s=m}^{t} \mathbb{P}\left(\frac{1}{s} \left| \sum_{i=1}^{s} \varepsilon_i \right| > \lambda \right) && (29)\\
&\leq 2 \sum_{s=m}^{t} \exp\left(-\frac{s\lambda^2}{2\sigma^2} \right) && (30)\\
&\leq 2 \int_{m}^{t+1} \exp\left(-\frac{s\lambda^2}{2\sigma^2} \right) ds && (31)\\
&\leq \frac{4\sigma^2}{\lambda^2} \exp\left(-\frac{m\lambda^2}{2\sigma^2} \right) && (32)
\end{aligned}$$

Now set

$$m = \frac{2\sigma^2(\log(T) - \log(\lambda))}{\lambda^2}, \tag{33}$$

and obtain

$$\mathbb{P}(A_m(i,\lambda)) \leq \frac{4\sigma^2}{\lambda^2} \exp\left(-\frac{m\lambda^2}{2\sigma^2}\right) = \frac{4\sigma^2}{\lambda T} \tag{34}$$

Bounding B_m From the above, we know that the estimation error concentrates well after observing m selections. Now we show that with high probability, it does not take too long to wait for those m selections. This follows from the probability of selection of an item at every time t is lower bounded by $p(1-p)^K$. In particular, for $M > m$, we note that the probability that we have not observed m selections is lower bounded by a Binomial random variable $Z \sim B(M, p(1-p)^K)$ since preferences are independent between agents. Consider

$$M = \frac{2\alpha m}{p(1-p)^K} = \frac{4\alpha\sigma^2(\log(T) - \log(\lambda))}{\lambda^2 p(1-p)^K} \tag{35}$$

where

$$\alpha = \max\left(1, \frac{2\lambda^2}{\sigma^2}\right). \tag{36}$$

First we note that in this case,

$$\mathbb{E}(Z) = 2\alpha m \geq 2m \tag{37}$$

and thus

$$\begin{aligned}
\mathbb{P}(B_m) &\leq \mathbb{P}(Z \leq m) && (38)\\
&\leq \mathbb{P}\left(Z \leq \frac{1}{2}\mathbb{E}(Z)\right) && (39)\\
&\leq \exp\left(-\frac{\mathbb{E}(Z)}{8}\right) && (40)\\
&\leq \exp\left(-\frac{\alpha\sigma^2(\log(T) - \log(\lambda))}{2\lambda^2}\right) && (41)\\
&\leq \exp\left(-(\log(T) - \log(\lambda))\right) && (42)\\
&= \frac{\lambda}{T} && (43)
\end{aligned}$$

where third inequality is a standard Chernoff bound and the second to last step follows from the condition on α.

Plugging these bounds on $A_m(i,\lambda)$ and $B_m(i,M)$ in to our bound for regret (27), we obtain

$$\mathbb{E}[\text{regret}(T)] \leq 2\left(\frac{4\sigma^2}{\lambda} + \lambda\right)K + \frac{8\alpha\sigma^2 K(\log(T) - \log(\lambda))}{\lambda^2 p(1-p)^K} + \lambda K T \mathbb{I}_{\lambda > \frac{\Delta_{\min}}{2}} \tag{44}$$

and thus if we set $\lambda = \frac{\Delta_{\min}}{2}$, we find

$$\mathbb{E}[\text{regret}(T)] \leq \left(\frac{16\sigma^2}{\Delta_{\min}} + \Delta_{\min}\right)K + \frac{32\alpha\sigma^2 K(\log(T) - \log(\Delta_{\min}) + \log(2))}{\Delta_{\min}^2 p(1-p)^K} \tag{45}$$

as desired. □